

ABSTRACT OF THE DISCLOSURE

A method and apparatus for speech recognition using latent semantic adaptation is described herein. According to one aspect of the present invention, a method for recognizing speech comprises using latent semantic analysis (LSA) to generate an LSA space for a collection of documents and to continually adapt the LSA space with new documents as they become available. Adaptation of the LSA space is optimally two-sided, taking into account the new words in the new documents. Alternatively, adaptation is one-sided, taking into account the new documents but discarding any new words appearing in those documents.

04860.P2638